Scalable Spectrum Allocation for Large Networks Based on Sparse Optimization

Binnan Zhuang
Modem R&D Lab
Samsung Semiconductor, Inc.
San Diego, CA

Dongning Guo, Ermin Wei, and Michael L. Honig
Department of Electrical Engineering and Computer Science
Northwestern University
Evanston, IL

arXiv:1702.05679v1 [cs.IT] 19 Feb 2017

***Abstract*—Joint allocation of spectrum and user association is considered for a large cellular network. The objective is to optimize a network utility function such as average delay given traffic statistics collected over a slow timescale. A key challenge is scalability: given n Access Points (APs), there are $O(2^n)$ ways in which the APs can share the spectrum. The number of variables is reduced from $O(2^n)$ to $O(nk)$, where k is the number of users, by optimizing over local overlapping neighborhoods, defined by interference conditions, and by exploiting the existence of sparse solutions in which the spectrum is divided into $k+1$ segments. We reformulate the problem by optimizing the assignment of subsets of active APs to those segments. An ℓ_0 constraint enforces a one-to-one mapping of subsets to spectrum, and an iterative (reweighted ℓ_1) algorithm is used to find an approximate solution. Numerical results for a network with 100 APs serving several hundred users show the proposed method achieves a substantial increase in total throughput relative to benchmark schemes.**

I. Introduction

Heterogeneous cellular networks with dense deployment of access points (APs) are anticipated to be a major component of 5G networks. A challenge with such dense deployments is interference management. Coordinated radio resource allocation across multiple cells is one approach for mitigating inter-cell interference. That includes joint scheduling across multiple cells over a fast timescale [1]–[3] as well as the assignment of resources across cells over a slower timescale [4], [5]. Whereas the former approach requires instantaneous knowledge of channel gains, the latter approach relies on statistical knowledge of interference.

We consider the joint allocation of spectrum and user association in a large network with many APs. The objective is to optimize a network utility, such as average delay, given traffic statistics over a geographic region that change slowly relative to channel fading. Our approach builds on our prior work [4], [6] in which for a network of n APs and k mobiles (or User Equipments (UEs)), the spectrum is partitioned into 2^n *patterns*, corresponding to all possible subsets of active APs. The problem is to optimize the widths of spectrum segments, associated with the different patterns, along with the user association under each pattern. This has been shown to provide significant performance improvement in throughput enhancement and delay reduction [4]–[7].

The original convex problem formulation in [4] is not useful for large networks because the number of patterns grows as $O(2^n)$. In prior work [6], we have reduced the number of variables to $O(n)$ by recognizing that each link rate depends only on *local* patterns of active APs. The problem can then be redefined over sets of overlapping *interference neighborhoods*, associated with those local patterns. Each AP has its own interference neighborhood, which captures the interference from nearby APs. The challenge with the approach in [6] is to ensure that the spectrum assigned to each particular AP is consistent across the neighborhoods to which it belongs. To accomplish that, the spectrum is discretized and a coloring algorithm is proposed to ensure that the local patterns of active APs are globally consistent.

Here we take a different approach to addressing the scalability problem. This is based on the fact that the solution is *sparse*, meaning that at most $k+1$ out of the 2^n possible patterns appear in the optimal allocation for a general network utility function. Hence we reformulate the problem by dividing the spectrum into $k+1$ segments, rather than 2^n, and attempt to identify the pattern that should be associated with each segment. This effectively reverses the approach in [6], which attempts to assign a segment of spectrum to each pattern. In this reformulation, we initially assume that any combination of patterns can be assigned to each of the $k+1$ segments. This problem is a convex relaxation of our original problem. The one-to-one mapping of spectrum segments to patterns is then enforced with an ℓ_0 (cardinality) constraint. An algorithm for finding an approximate solution to this problem is presented based on a reweighted ℓ_1 approximation of this constraint [8].

The approach to scalability presented here has the advantage of eliminating the combinatorial coloring problem that arises in [6]. Instead, a new ℓ_0 constraint is introduced. Although this does not simplify the original problem, it helps in finding an approximate solution, since the reweighted ℓ_1 approximations for the ℓ_0 constraints are known to perform well. Numerical results indicate that this method generally gives better performance for a fixed computational complexity than the method in [6].

II. System Model

As in [6], we consider a network containing the set of APs $\mathcal{N} = \{1, \cdots, n\}$ and the set of UEs $\mathcal{K} = \{1, \cdots, k\}$. Each "UE" is actually associated with a particular location, and could refer to a group of nearby mobiles. The n APs share W Hz of spectrum. For convenience, we normalize

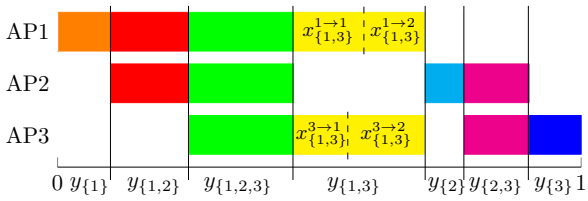

Fig. 1. An example off spectrum allocation among 3 APs.

$W = 1$. Each AP can transmit on any part(s) of the spectrum, which is homogeneous (has the same distance attenuation). APs sharing the same spectrum interfere at UEs within range of both transmitters. All transmissions are assumed to be omni-directional, although a similar set of problems can be reformulated with directional transmissions.

We define a subset of active (transmitting) APs $A \subset \mathcal{N}$ as a *pattern*. There are 2^n patterns, and a particular allocation of spectrum maps each pattern to a slice of spectrum. We denote an allocation as $\{y_A\}_{A\subset\mathcal{N}}$, where y_A denotes the amount of bandwidth assigned to pattern A. An example with three APs is depicted in Fig. 1. AP 1 owns $\{1\}$ exclusively; shares $\{1,2\}$ with AP 2; shares $\{1,3\}$ with AP 3; and shares $\{1,2,3\}$ with both AP 2 and AP 3.

User association is determined by how the spectrum assigned to a particular AP is allocated to different UEs. Specifically, denote the fraction of total bandwidth used by AP i to serve UE j under pattern A as $x_A^{i\to j}$, for $i \in A$. UE j is then assigned to AP i if $x_A^{i\to j} > 0$ for some $A \subset \mathcal{N}$. Since the total bandwidth assigned to pattern A is y_A, we have

$$\sum_{j\in\mathcal{K}} x_A^{i\to j} \le y_A, \ \forall A \subset \mathcal{N}, \ i \in A. \tag{1}$$

The total bandwidth allocated to all patterns is then

$$\sum_{A\subset\mathcal{N}} y_A = 1. \tag{2}$$

Let $s_A^{i\to j}$ denote the spectral efficiency of the link from AP i to UE j under pattern A. This is measured over a *slow timescale*, and is therefore an average over short-term fading. The value of $s_A^{i\to j}$ is therefore determined by the distance between AP i and UE j, shadowing, and similar long-term characteristics of the interference links from APs in A to UE location j. For concreteness we assume

$$s_A^{i\to j} = \frac{W}{\tau}\mathbb{1}_{\{i\in A\}}\log_2\left(1 + \frac{p_i g^{i\to j}}{\sum_{i'\in A\setminus\{i\}} p_{i'} g^{i'\to j} + n_j}\right) \tag{3}$$

where $\mathbb{1}_{\{i\in A\}} = 1$ if $i \in A$ and 0 otherwise, p_i is the transmit power spectral density (PSD) at AP i, $g^{i\to j}$ is the power gain of link $i \to j$, and n_j is the noise PSD at UE j. We assume fixed, flat transmit PSDs over the slow timescale considered. The factor W/τ, where τ is the average packet length (bits), gives the units in packets/sec. The link gain $g^{i\to j}$ includes pathloss and shadowing effects. Clearly, $s_A^{i\to j} = 0$ if $i \notin A$, i.e., AP i does not transmit on pattern A. In practice, the spectral efficiencies can be measured as time-averaged channel gains. The total rate received by UE j is therefore

$$r^j = \sum_{A\subset\mathcal{N}}\sum_{i\in A} s_A^{i\to j} x_A^{i\to j}, \ \forall j \in \mathcal{K}. \tag{4}$$

III. Problem Formulation with Global Patterns

The problem is to maximize a network utility function over the spectrum allocation and user association designated by $\boldsymbol{x} = \left(x_A^{i\to j}\right)_{A\subset\mathcal{N}, i\in A, j\in\mathcal{K}}$, $\boldsymbol{y} = (y_A)_{A\subset\mathcal{N}}$:

$$\underset{\boldsymbol{r},\,\boldsymbol{x},\,\boldsymbol{y}}{\text{maximize}} \ u(r^1,\cdots,r^k) \tag{P0a}$$

$$\text{subject to } x_A^{i\to j} \ge 0, \quad \forall A \subset \mathcal{N}, \ \forall i \in A, \ \forall j \in \mathcal{K} \tag{P0b}$$

and constraints (4), (1), and (2), where the network utility function u depends on the service rates to all UEs. The optimization problem is convex if u is concave in $\boldsymbol{r} = [r^1,\cdots,r^k]$. As in [6], we will take u to be the average packet delay, given by

$$u(r^1,\cdots,r^k) = -\sum_j \frac{\lambda^j}{(r^j-\lambda^j)^+} \tag{5}$$

where $\frac{1}{(x)^+}$ equals $\frac{1}{x}$ if $x > 0$ and $+\infty$ otherwise, and λ_j is the Poisson packet arrival rate for UE j. This assumes exponential packet lengths and backlogged interference [4].

Solving P0 becomes prohibitively expensive as the network size grows, due to the inherit complexity from the 2^n global patterns. However, a key property of solution(s) to (P0) is that at least one is *sparse*, i.e., contains at most $k+1$ patterns.

Proposition 1: ([9]) P0 has a solution that divides the spectrum into at most $k+1$ segments, i.e.,

$$|\{A \subset \mathcal{N} \mid y_A > 0\}| \le k+1. \tag{6}$$

Furthermore, if u is element-wise nondecreasing in $\boldsymbol{r}$, $k+1$ in (6) is reduced to k.
Determining which $k+1$ active patterns appear in a solution is then the key challenge in solving P0.

IV. Reformulation With Sparsity Constraints

A. Local Neighborhoods

We first reduce the number of variables in (P0) from $O(2^n)$ to $O(n)$ by reformulating (P0) over *interference neighborhoods* based on local patterns [6]. Due to pathloss, we assume interference vanishes beyond a certain distance. Let $L \subset \mathcal{N}\times\mathcal{K}$ denote the set of links with nonzero gains. Hence each UE only receives power from APs within its neighborhood:

$$\mathcal{A}_j = \{i | (i \to j) \in L\}. \tag{7}$$

From the AP side, each AP can only transmit to a collection of UEs in a AP neighborhood with positive rates:

$$\mathcal{U}_i = \{j | (i \to j) \in L\}. \tag{8}$$

We define the interference neighborhood for each AP i as:

$$\mathcal{N}_i = \cup_{j\in\mathcal{U}_i}\mathcal{A}_j, \tag{9}$$

i.e., $\mathcal{N}_i$ includes AP i and all APs that interfere with it.

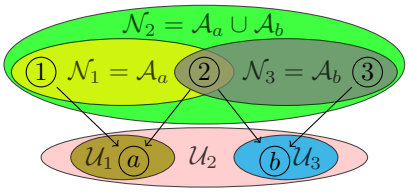

Fig. 2. Neighborhoods in the case of three APs and two UEs.

An example with three APs (denoted by $\{1,2,3\}$) and 2 UEs (denoted by $\{a,b\}$) is shown in Fig 2. The set of nonzero links are $L = \{1 \to a,\ 2 \to a,\ 2 \to b,\ 3 \to b\}$. The AP neighborhoods are $\mathcal{U}_1 = \{a\}$, $\mathcal{U}_2 = \{a,b\}$, and $\mathcal{U}_3 = \{b\}$; and the UE neighborhoods are $\mathcal{A}_a = \{1,2\}$ and $\mathcal{A}_b = \{2,3\}$. AP 1's interference neighborhood is $\mathcal{N}_1 = \{1,2\}$, as AP 2 interferes with it at UE a. AP 3's interference neighborhood is $\mathcal{N}_3 = \{2,3\}$, as AP 2 interferes with it at UE b. AP 2's interference neighborhood is $\mathcal{N}_2 = \{1,2,3\}$, since AP 1 and AP 3 interfere with AP 2 at UE a and UE b, respectively.

The spectral efficiency of link $i \to j$ only depends on the local patterns in $\mathcal{A}_j$, according to the definition of UE neighborhoods in (7). Therefore, for any link $i \to j$, the spectral efficiency under a global pattern A is equivalent to its intersection with UE j's neighborhood $\mathcal{A}_j$:

$$s_A^{i\to j} = s_{A\cap\mathcal{A}_j}^{i\to j},\ \forall j \in \mathcal{K},\ \forall i \in \mathcal{A}_j,\ A \subset \mathcal{N}. \tag{10}$$

We next express each rate in terms of local patterns. We define bandwidth allocation variables within a local interference neighborhood $\mathcal{N}_i$ as:

$$z_B^{i\to j} = \sum_{A\subset\mathcal{N}: A\cap\mathcal{N}_i = B} x_A^{i\to j},\ i \in \mathcal{N},\ j \in \mathcal{U}_i,\ \forall B \subset \mathcal{N}_i. \tag{11}$$

Because all global patterns sharing the same overlap with $\mathcal{N}_i$ contribute to the same local pattern of AP i, $z_B^{i\to j}$ represents the bandwidth allocated to link $i \to j$ under local pattern B. Hence from AP i's perspective, the bandwidth assigned to a local pattern B must be the sum of the bandwidths assigned to all global patterns containing B. The service rate of link $i \to j$, as defined in (4), can then be calculated as [6]:

$$r^j = \sum_{i\in\mathcal{A}_j} \sum_{B\subset\mathcal{N}_i} s_{B\cap\mathcal{A}_j}^{i\to j} z_B^{i\to j} \tag{12}$$

where we have used (10) and (11), and the fact that only APs in $\mathcal{A}_j$ transmit to UE j with positive rate.

B. Sparse Optimization

Motivated by Proposition 1, we reformulate (P0) by dividing the spectrum into $k+1$ segments, and seek to assign a single pattern to each segment. Proposition 1 implies that by optimizing this assignment we obtain a solution to P0. Let h_l denote the bandwidth of the lth segment (to be optimized). We now associate a set of variables z (for local neighborhoods) and y *for each segment* l. We can therefore rewrite the rate in (12) as

$$r^j = \sum_{l=1}^{k+1} \sum_{i\in\mathcal{A}_j} \sum_{B\subset\mathcal{N}_i} s_{B\cap\mathcal{A}_j}^{i\to j} z_{B,l}^{i\to j},\ \forall j \in \mathcal{K} \tag{13}$$

where the first sum is over the $k+1$ spectrum segments. Apart from the addition of segment index, the z variable is the same local bandwidth allocation variable defined in (11). The amount of bandwidth assigned to local pattern B in AP i's interference neighborhood $\mathcal{N}_i$ within segment l is

$$y_{B,l}^i = \sum_{j\in\mathcal{U}_i} z_{B,l}^{i\to j},\ \forall i \in \mathcal{N},\ \forall B \subset \mathcal{N}_i. \tag{14}$$

The total amount of spectrum assigned to $\mathcal{N}_i$ satisfies

$$\sum_{B\subset\mathcal{N}_i} y_{B,l}^i \le h_l,\ \forall i. \tag{15}$$

We also introduce the following *consistency constraint* to ensure that the amount of spectrum allocated to an AP is consistent across any two neighborhoods $\mathcal{N}_i$ and $\mathcal{N}_m$ that contain it [6]. That is, for any nonempty $C \subset \mathcal{N}_i \cap \mathcal{N}_m$,

$$\sum_{B\subset\mathcal{N}_i: B\cap\mathcal{N}_m = C} y_{B,l}^i = \sum_{B\subset\mathcal{N}_m: B\cap\mathcal{N}_i = C} y_{B,l}^m. \tag{16}$$

See, for example, $\mathcal{N}_1$ and $\mathcal{N}_3$ in Fig. 2. In interference cluster $\mathcal{N}_1$, AP 2 transmits under pattern $\{2\}$ and $\{1,2\}$; in interference cluster $\mathcal{N}_3$, AP 2 transmits under pattern $\{2\}$ and $\{2,3\}$. The total bandwidth used by AP 2 must be consistent across neighborhoods so that $y_{\{2\},l}^1 + y_{\{1,2\},l}^1 = y_{\{2\},l}^3 + y_{\{2,3\},l}^3$, where $\{2\}$ is the overlapping pattern. This example can be extended to a *set* of APs, which are members of two interference neighborhoods, giving (16).

To ensure a one-to-one mapping of patterns to the $k+1$ segments, we add the ℓ_0-norm constraint

$$\sum_{B\subset\mathcal{N}_i} |y_{B,l}^i|_0 \le 1,\ \forall i \in \mathcal{N},\ l = 1,\cdots,k+1 \tag{17}$$

where $|x|_0 = 1$ if $x \neq 0$, and $|x|_0 = 0$ if $x = 0$. That is, we constrain each AP to use at most one active pattern in each segment.

We can now reformulate P0 in terms of the local interference variables $\boldsymbol{z}$ across the $k+1$ spectrum segments:

$$\underset{\boldsymbol{r},\,\boldsymbol{y},\,\boldsymbol{z}}{\text{maximize}}\ u(r^1,\cdots,r^k) \tag{P1a}$$

$$\text{subject to}\ z_{B,l}^{i\to j} \ge 0,\quad \forall (i\to j) \in L,\ \forall B \subset \mathcal{N}_i, \tag{P1b}$$

$$\sum_{l=1,\cdots,k+1} h_l = 1,\quad h_l \ge 0, \tag{P1c}$$

and constraints (13)-(17) for each $l = 1,\cdots,k+1$.

Theorem 1: P0 and P1 are equivalent (have the same solutions) given the local neighborhood definitions (7)-(9).

The proof is omitted due to limited space. Hence a solution to P1 always satisfies Proposition 1. This reformulation leads to a computationally efficient approximation algorithm.

V. Iterative Approximation

The number of variables is reduced from $O(2^n)$ in P0 to $O(nk)$ in P1. However, the ℓ_0 norm constraint makes the problem non-convex and difficult to solve. We apply the reweighted ℓ_1 approach, described in [8], to approximate the ℓ_0 constraint. Specifically, the ℓ_0 norm is approximated by the

weighted ℓ_1 norm, $\sum_i w_i|x_i|$, where w_i is iteratively adapted, and taken to be inversely proportional to $|x_i|$ computed at the preceding iteration. This has the effect of suppressing small nonzero entries with large weights.[1]

The ℓ_1 approximation cannot be directly applied to P1 since the ℓ_0 norm constraints are coupled through (17) and (P1c). Hence we present an iterative algorithm based on the ℓ_1 reweighted heuristic, where the weights depend on both the ℓ_1 norm and the bandwidth of each segment h_l. In each iteration of the algorithm, shown in Algorithm 1, we solve P1, but with the ℓ_0 norm constraint (17) replaced by the weighted sum

$$\sum_{B\subset\mathcal{N}_i} w_{B,l}^i y_{B,l}^i \leq 1. \tag{18}$$

In Algorithm 1, we refer to this reweighted ℓ_1 version of P1 as P2.

Algorithm 1 shows a random initialization of the weights to introduce asymmetry in the first iteration. Otherwise, e.g., if all $w_{B,l}^i$'s are initialized to a constant, the solution stays symmetric over all segments, because the constraints and objectives are identical for each segment. A symmetric solution, i.e., $y_{B,1}^i = y_{B,2}^i =, \cdots, = y_{B,k}^i,\ \forall i \in \mathcal{N},\ B \subset \mathcal{N}_i$, is generally not a solution to the original problem. In each iteration, we solve P2 with the current weights $\boldsymbol{w}$ to obtain the current $\boldsymbol{x}$, $\boldsymbol{y}$, $\boldsymbol{z}$ and $\boldsymbol{h} = [h_1, \cdots, h_k]$. Then the weights are updated as shown. The iterations terminate when the variables converge or the maximum number of iterations t_{max} is reached.

The weight update in Algorithm 1 is obtained by approximating the ℓ_0 norm with $\log(x+\epsilon)$ for small ϵ (see [8] and the references therein). Adding ϵ in the denominator allows small components to be propagated to the next iteration. Here the goal is to obtain a *single* nonzero $y_{B,l}^i = h_l,\ B \in \mathcal{N}_i$, so we take $\epsilon = \alpha h_l$. This is different from the fixed ϵ proposed in [8] in that αh_l changes with h_l in each iteration. That is, Algorithm 1 simultaneously searches for the optimal reuse pattern and its associated bandwidth.

It is possible that Algorithm 1 produces multiple reuse patterns for a segment at termination. In those cases AP i chooses the dominant pattern $B_l^i = \arg\max_{B\in\mathcal{N}_i} y_{B,l}^i$. The global reuse pattern assigned to segment l is then given by $A_l^* = \cup_{i:i\in B_l^i}\{i\}$. The spectral efficiencies for each segment are determined by the corresponding global reuse pattern (set of interfering APs): $\bar{s}_l^{i\to j} = s_{A_l^*\cap\mathcal{A}_j}^{i\to j}$. Given the assignment of patterns to each spectrum segment, the widths of the segments, $\boldsymbol{h} = (h_l)_{l=1,\cdots,k+1}$ along with the assignment of spectrum to mobiles, $\bar{\boldsymbol{x}} = (\bar{x}_l^{i\to j})_{j\in\mathcal{K},i\in\mathcal{A}_j,l=1,\cdots,k+1}$, can then be re-optimized by solving the convex problem:

$$\underset{\bar{\boldsymbol{x}},\,\boldsymbol{h}}{\text{maximize}}\ u(r_1,\cdots,r_k) \tag{P3a}$$

$$\text{subject to } r_j = \sum_{l=1}^{k}\sum_{i\in\mathcal{A}_j} \bar{s}_l^{i\to j}\bar{x}_l^{i\to j}, \quad \forall j \in \mathcal{K} \tag{P3b}$$

$$\sum_{j=1}^{k} \bar{x}_l^{i\to j} \leq h_l,\ \forall j \in \mathcal{K}, i \in \mathcal{A}_j \tag{P3c}$$

$$\bar{x}_l^{i\to j} \geq 0 \text{ and (P1c) for } l = 1,\cdots,k+1.$$

[1]This algorithm has also been used in [9] to solve an AP activation problem.

Algorithm 1 Iterative re-weighted ℓ_1 approximation.

INPUT: $(s_C^{i\to j})_{j\in\mathcal{K},i\in\mathcal{A}_j,C\subset\mathcal{A}_j}$, and $(\lambda_j)_{j\in\mathcal{K}}$.

OUTPUT: The widths of $k+1$ segments $(h_l)_{l=1,\cdots,k+1}$, the $k+1$ active patterns $(A_l^*)_{l=1,\cdots,k+1}$, and the spectrum allocated to link $i \to j$ on segment l, $(\bar{x}_l^{i\to j})_{j\in\mathcal{K},i\in\mathcal{A}_j,l=1,\cdots,k+1}$

Initialization: Choose $w_{B,l}^i$'s randomly in (0,1).

Set iteration counter $t = 0$.

while Variables have not converged and $t < t_{max}$ **do**

1. Solve P2 with the current weights $\boldsymbol{w}$.
2. Update $w_{B,l}^i = \frac{1}{y_{B,l}^i+\alpha h_l}$.
3. $t = t + 1$.

end while

Post Processing: Determine the reuse patterns $\{A_l^*\}$ and spectral efficiencies $\{\bar{s}_l^{i\to j}\}$ across segments; solve P3.

VI. Numerical Results

The simulation results assume one macro AP is located at the center of the area, and the remaining $n-1$ pico APs are randomly dropped around it. The k UEs are placed on a rectangular lattice. Link gains include both pathloss and shadowing. Additional simulation parameters are shown in the footnote.[2]

A. Small Network

We first compare solutions to P0 and P1 for a small network with $n = 10$ and $k = 32$. Since the number of variables is relatively small, we solve both versions of P0 with and without the local neighborhood approximation using a standard convex optimization solver. The local neighborhoods are constructed by including the strongest four APs for each UE. The solution to P1 is obtained using Algorithm 1. We compare those with full spectrum reuse where each user is assigned to the AP with the strongest signal (maxRSRP association), and also the optimal orthogonal allocation,[3] i.e., only $\{y_{\{i\}}\}_{i\in\mathcal{N}}$ are active.

Fig. 3 shows delay versus traffic arrival rate for all schemes. The end of each curve represents the maximum arrival rate the scheme can support. The curves obtained by solving P0 with and without local neighborhood approximation are very close, which indicates considering the four strongest interferers is enough in such a small network. The solution to P1 incurs slightly larger delay than the solution to P0. The jointly optimized spectrum allocations and user associations achieve substantial delay reduction as well as eight times throughput compared to full frequency reuse with maxRSRP association.

The optimized spectrum allocation and user associations obtained solving P1 are depicted in Fig. 4. The macro and pico APs are represented by the bigger and smaller towers; each

[2]The pathloss exponent is 3, standard deviation of shadow fading is 3, macro-transmit PSD is 5 μW/Hz, pico transmit PSD is 1 μW/Hz, noise PSD is 10^{-7} μW/Hz, total bandwidth is 20 MHz, and average packet length is 1 Mb.

[3]Both spectrum allocation and user association are optimized assuming each AP exclusively occupies a fraction of the spectrum.

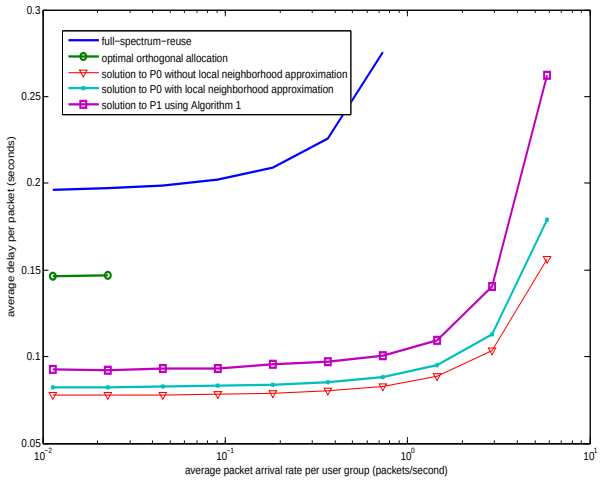
full−spectrum−reuse
optimal orthogonal allocation
solution to P0 without local neighborhood approximation
solution to P0 with local neighborhood approximation
solution to P1 using Algorithm 1
average delay per packet (seconds)
average packet arrival rate per user group (packets/second)

Fig. 3. Delay versus traffic arrival rate curves for a small network with $n = 10$ and $k = 23$.

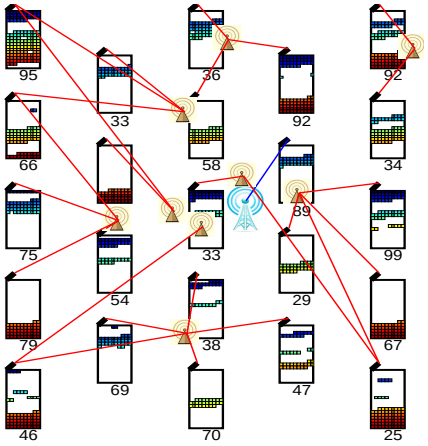

Fig. 4. Example spectrum allocations and user associations from (P1).

handset represents a UE. Each solid line shows a connection between the corresponding AP and UE. The grid for each UE shows the spectrum used to serve that UE. The traffic arrival rate (in (0, 100)) for each UE is shown under each grid. The user association in Fig. 4 is close but not identical to that obtained by solving P0 (not shown) due to the approximation in Algorithm 1, which explains the performance difference shown in Fig. 3.

B. Large Network

Fig. 5 compares the performance of different allocation schemes in a large network with $n = 100$ APs and $k = 200$ UEs. To facilitate the simulations, we reduce the size of each UE neighborhood from four to three, i.e., each UE can only be served by the three strongest APs. We compare full-spectrum-reuse with maxRSRP association, full-spectrum-reuse with optimized associations and the solution to P1 with 50 segments. The solution to P1 achieves 1.5 times the throughput of the full-spectrum-reuse with optimized association. The solutions to P1 use no more than 23 active patterns (< 50 available segments). The throughput gain achieved by the solution to P1

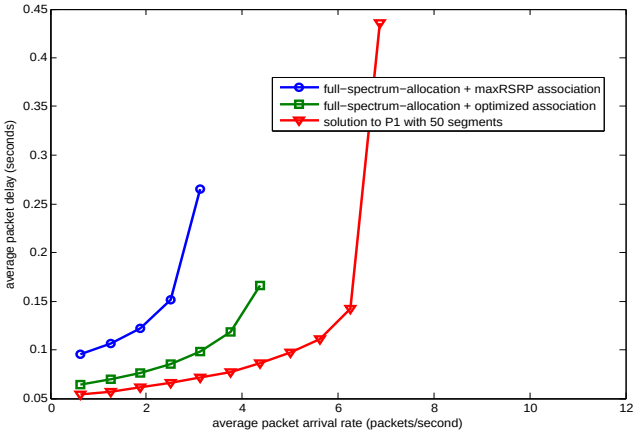
full−spectrum−allocation + maxRSRP association
full−spectrum−allocation + optimized association
solution to P1 with 50 segments
average packet delay (seconds)
average packet arrival rate (packets/second)

Fig. 5. Delay versus traffic arrival rate curves for a large network with $n = 100$ and $k = 195$.

in this large network is smaller than that in Section VI-A. This is mainly because we only consider the three strongest interferers, which compromises the benefits from interference management.

VII. Conclusion

An approach to joint allocation of spectrum with user association has been presented which exploits the sparsity of the optimal solution. The proposed algorithm has been observed to achieve near-optimal performance for small to medium-size networks for which the optimal solution can be computed, and provides substantial gains relative to full frequency reuse. Although not considered here, the formulation can be extended to accommodate spatial selectivity and different power levels. The performance-complexity tradeoff for such extensions is left for future work.

References

[1] W. Yu, T. Kwon, and C. Shin, "Multicell coordination via joint scheduling, beamforming, and power spectrum adaptation," *IEEE Trans. Wireless Commun.*, vol. 12, no. 7, pp. 1–14, 2013.

[2] F. Wang, L. Song, Z. Han, Q. Zhao, and X. Wang, "Joint scheduling and resource allocation for device-to-device underlay communication," in *Proc. Conf. Wireless Comm. and Networking*, pp. 134–139, IEEE, 2013.

[3] M. Hong and Z.-Q. Luo, "Distributed linear precoder optimization and base station selection for an uplink heterogeneous network," *IEEE Trans. Signal Process.*, vol. 61, pp. 3214–3228, June 2013.

[4] B. Zhuang, D. Guo, and M. L. Honig, "Traffic-driven spectrum allocation in heterogeneous networks," *IEEE J. Sel. Areas Commun.*, vol. PP, no. 99, pp. 1–1, 2015.

[5] Q. Kuang, "Joint user association and reuse pattern selection in heterogeneous networks," in *2014 11th International Symposium on Wireless Communications Systems (ISWCS)*, pp. 401–405, IEEE, 2014.

[6] B. Zhuang, D. Guo, E. Wei, and M. L. Honig, "Scalable spectrum allocation and user association in networks with many small cells," *https://arxiv.org/abs/1701.03247*, submitted to *IEEE Trans. Commun.*.

[7] Q. Kuang and W. Utschick, "Energy management in heterogeneous networks with cell activation, user association, and interference coordination," *IEEE Trans. Wireless Commun.*, vol. 15, June 2016.

[8] E. J. Candes, M. B. Wakin, and S. P. Boyd, "Enhancing sparsity by reweighted l_1 minimization," *Journal of Fourier analysis and applications*, vol. 14, no. 5-6, pp. 877–905, 2008.

[9] B. Zhuang, D. Guo, and M. L. Honig, "Energy-efficient cell activation, user association, and spectrum allocation in heterogeneous networks," *IEEE J. Sel. Areas Commun.*, vol. PP, no. 99, pp. 1–1, 2016.